Effective Defect Detection Using Instance Segmentation for NDI

Ashiqur Rahman[1], Venkata Devesh Reddy Seethi[1], Austin Yunker[2], Zachary Kral[3], Rajkumar Kettimuthu[2], and Hamed Alhoori[1]

[1]Department of Computer Science, Northern Illinois University, DeKalb, IL 60115 USA
[2]Argonne National Laboratory, Lemont, IL, USA
[3]Spirit AeroSystems, Wichita, KS, USA
ashiqur.r@niu.edu

Abstract

Ultrasonic testing is a common Non-Destructive Inspection (NDI) method used in aerospace manufacturing. However, the complexity and size of the ultrasonic scans make it challenging to identify defects through visual inspection or machine learning models. Using computer vision techniques to identify defects from ultrasonic scans is an evolving research area. In this study, we used instance segmentation to identify the presence of defects in the ultrasonic scan images of composite panels that are representative of real components manufactured in aerospace. We used two models based on Mask-RCNN (Detectron 2) and YOLO 11 respectively. Additionally, we implemented a simple statistical pre-processing technique that reduces the burden of requiring custom-tailored pre-processing techniques. Our study demonstrates the feasibility and effectiveness of using instance segmentation in the NDI pipeline by significantly reducing data pre-processing time, inspection time, and overall costs.

Introduction

Non-Destructive Inspection (NDI) is inspecting an object without damaging it. This process is essential in manufacturing, construction, defect evaluation, art restoration, infrastructure safety, and many more (Honarvar and Varvani-Farahani 2020; Ould Naffa et al. 2002; Dwivedi, Vishwakarma, and Soni 2018; Memmolo et al. 2015; Brosnan and Sun 2004). With the advent of artificial intelligence (AI) and machine learning (ML), automation of the NDI process has become at the forefront of research (Gardner et al. 2020; Lin et al. 2023). In this study, we worked with ultrasonic scans of thin-walled composite materials, used in the aerospace manufacturing industry to identify defects in the manufactured product. A trained inspector inspects the scan using proprietary software to identify any defects in the scans. In our work, we want to augment the process by suggesting areas of interest for the inspector to look for defects. With the help of ML, we can reduce the inspection time and inspector fatigue and increase the safety of the manufactured parts.

The process of using ML to identify defects from ultrasonic scans is of great interest to the NDI research community. Several studies suggested methods of identifying defects using ML (Baumgartl et al. 2020; Yang et al. 2022). However, the lack of high-quality training data and high training time are constraints to the progress of this field. While researchers agree that ML is the direction to choose for the future of NDI (Jodhani et al. 2023; Sun, Ramuhalli, and Jacob 2023), there are still many improvements to make before it becomes the primary choice (Hassani and Dackermann 2023). In the manufacturing industry, the high cost of deployment and maintenance is also a prohibitive factor. So, for ML-based approaches to have a place in the NDI pipeline, the training and maintenance of the models need to be cheaper, and the method should have very high accuracy to ensure the safety of the critical parts. In this work, we propose image segmentation models as a feasible approach for ML-based NDI systems. While image segmentation is used by many researchers as a tool to pre-process the data, or as a part of the defect detection models, using image segmentation as a viable defect detection approach needs further research.

Object detection using image segmentation models has recently become a rising research area. The object detection models classify images at the pixel level to identify the presence of an object (Hafiz and Bhat 2020). In this study, we used trained two different models - Detectron 2 based on Mask-RCNN (Wu et al. 2019a) and "You Only Look Once" (YOLO) 11 (Redmon et al. 2016; Khanam and Hussain 2024) to detect defects in the ultrasonic scans and compared the results to propose this as a viable approach for NDI.

Related Works

NDI. NDI is an essential part of the manufacturing pipeline that ensures the quality of manufactured products through inspection without destroying or breaking the products (Gholizadeh 2016; Honarvar and Varvani-Farahani 2020). This inspection process ensures the reliability of the product right from the start. Additionally, in safety-critical manufacturing domains, such as aerospace, NDI is necessary for the scheduled inspections that are conducted when the product is in use, to ensure there is no major damage occurring from wear and tear (Khedmatgozar Dolati et al. 2021).

NDI Techniques. NDI is an undeniable part of manufacturing, the food industry, structural health monitoring, inspecting artifacts, and many more. In the airline industry, the inspection of composite material is an essential part of NDI process. Studies such as Diamanti, Soutis, and Hodgkinson (2005); Gholizadeh (2016); Dwivedi, Vishwakarma, and Soni (2018); Honarvar and Varvani-Farahani (2020) showed different effective techniques of inspection, including but not limited to, Visual Testing, Ultrasonic Testing, Thermography, Radiographic Testing, Electromagnetic Testing, Acoustic Emission, Shearography Testing, and so on. With the advancement of technology over time, composite structures are preferred over metals for their lightweight and durability in the aerospace industry. Metal parts used eddy current and magnetic particle induction techniques for NDI (Lange 1994). However, these techniques do not work on composite structures like carbon fiber-reinforced polymer (CFPR) materials. Advances in NDI techniques addressed this in studies from Gupta et al. (2022); Gholizadeh (2016). Ultrasonic testing is the most common method used in the industry (Honarvar and Varvani-Farahani 2020). Despite its inability to detect very small defects and dependence on the experience of the inspectors (Gupta et al. 2022), ultrasonic testing, due to its ability to detect sub-surface defects and relatively low cost, has been the most popular method for the NDI process (Dwivedi, Vishwakarma, and Soni 2018; Wang et al. 2020). The ultrasonic-based NDI approach requires processing complex signals to interpret for the presence of flaws in a structure. Due to this complexity, the use of ML has become inevitable.

ML in NDI. Many research articles evaluated (Mishra, Bhatia, and Maity 2020) and discussed the importance, usefulness, and challenges (Gardner et al. 2020; Taheri, Gonzalez Bocanegra, and Taheri 2022; Lin et al. 2023) of ML in NDI.

Mishra, Bhatia, and Maity (2020) evaluated artificial neural networks (ANN), adaptive neuro-fuzzy inference systems (ANFIS), and support vector regression (SVR) models and found SVR most promising. While discussing the undeniable performance improvement of NDI with the help of ML. Gardner et al. (2020) also underscored the challenges of data availability, data quality, and complexity in training the ML models. Other methods like gradient boosting decision tree (GBDT) (Yang et al. 2022), principal component analysis (PCA), and support vector machine (SVM) (Ma, Tsuchikawa, and Inagaki 2020) are proposed for NDI techniques. However, appropriate hyper-parameter tuning and quality training data are critical for the better performance of ML (Sun, Ramuhalli, and Jacob 2023).

Image Segmentation. Image segmentation is an incremental research area that has evolved over decades and is rooted in classification research. The instance segmentation techniques, which have become highly effective in recent times, combine the detection of object bounding boxes and categorically detecting every image pixel (Hafiz and Bhat 2020). Image segmentation on ultrasonic images is a crucial technique in the computer-aided diagnosis (CAD) field, where researchers use semantic segmentation techniques to identify cancerous areas in the ultrasonic images (Su et al. 2011; Irfan et al. 2021). A similar technique can be used in ultrasonic scan images of other materials to identify defects. Several deep learning models have also been proposed by the research community like UPSNet, Panoptic-DeepLab, Detectron, and so on (Kirillov et al. 2019; Xiong et al. 2019; Cheng et al. 2020; Wu et al. 2019b). The available models use medical or real-world object datasets for training (Elharrouss et al. 2021). In aerospace, several works have employed semantic segmentation for identifying aircraft components such as engine, wing, and fuselage using DeepLab $v3$ (Thomas et al. 2024) and YOLO $v5$ (Xiang, Chang, and Ye 2024). Furthermore, computer vision was also used for identifying surface defects on aircraft from drone images using YOLO-FDD (Li, Wang, and Liu 2024). However, identifying defects from $3D$ ultrasonic images is challenging due to the images requiring additional pre-processing to transform into $2D$ images. Prakash et al. (2023) have used a software to first visualize the fuselage scan, apply complex feature processing techniques such as histogram of gradients, and train a KNN classifier to identify defects.

In this study, we aim to use simple pre-processing techniques on raw $3D$ ultrasonic scans and train instance segmentation models to identify defects. Furthermore, we do a comparative analysis of two popular segmentation frameworks such as Detectron 2 and YOLO 11, and evaluate the computational needs as well as their performance. Developing instance segmentation models that can achieve desirable performance with minimal pre-processing, enhances the scalability of the method as well as reduce the time and computational cost.

Data

We used a proprietary dataset containing scans of specimen panels, which are thin-walled composite structures representative of aerospace structures used in the industry. The panels consist of multiple plies of unidirectional carbon fiber polymer. In some of the specimens, an additional fiberglass ply was added on top of the layup. The specimen has a range of thicknesses to account for changes in actual structure. Teflon strips were placed throughout the layup at controlled locations to represent defects in the material. Per industry standards, the Teflon strips at varying depths and locations. The Teflon material has a sound impedance that can represent foreign inclusions or delaminations in the composite structure.

A single transducer system with $2.5MHz$ and $5MHz$ frequencies is used to scan the plates from both sides. The transducer was positioned perpendicular to the structure and sent ultrasonic waves through the material, and a receiver at the opposite of the object received the signals. From each point of scans, 512, 1024, and 2048 samples were collected for further analysis. Fluctuations in the amplitude of the received signals indicate the presence of defects. Differences in the signal attenuation can be observed from the visualization in Figure 1.

Each scan comprising the signal data has height and width dimensions of 258×368 and channels 512, 1024, or 2048, depending on sampling frequency as discussed above. Since

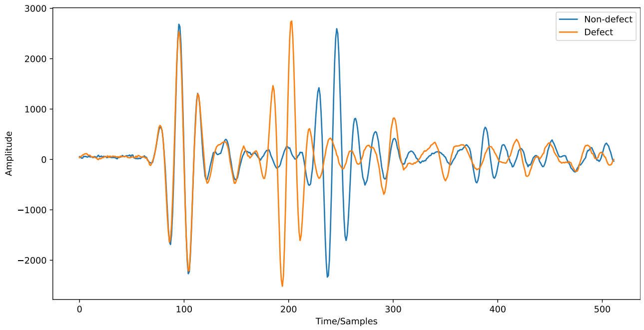

Figure 1: Sample of signals showing the difference between defect and non-defect areas

we use 2D segmentation models, we first convert these $3D$ scans to $2D$ data by taking the variance over the sampling frequency dimension. Then, we converted the $2D$ data into *NumPy* arrays and exported them as *PNG* images. For the purpose of training the models, we converted the defect annotations into COCO formatted (Lin et al. 2014) JSON files. The annotated JSON contained the bounding boxes for each defect which will be used for the training. To reduce the resource usage during training, we converted the images into grayscale. Using this process, we created 72 files for our study. We randomly selected 56 images for training, 8 for validation, and 8 for testing our models. Figure 2 (A) shows an example of the exported images used for training. Orange labels in Figure 2 (B) indicate the defect regions on the training image.

Method

We used two different models to identify the defects in the exported images. We used the instance segmentation approach, which is a more refined version of semantic segmentation, where each pixel of an image is classified, and a cluster of pixels is classified as part of a class (Hafiz and Bhat 2020). We used two different pre-trained algorithms, Detectron 2 based on Mask-RCNN (Wu et al. 2019a) and YOLO 11 (Jocher, Jing, and Chaurasia 2023), and fine-tuned them to train on our datasets.

Mask R-CNN (Detectron 2)

We used the instance segmentation feature of the Detectron 2 library, which utilizes the Mask R-CNN model. In this approach, the model first detects an object and creates a bounding box. Subsequently, it classifies the pixels inside that bounding box. To optimize model performance, we trained the model using various configurations and hyperparameter settings, such as number of epochs, batch size, images per batch, etc.

Training Configuration. To train the Detectron 2 model, we prepared the COCO dataset discussed above. The training was performed on a machine with $12GB$ GPU and $64GB$ RAM. To find the optimal model configuration, we conducted a series of training experiments using various fine-tuned hyperparameters, as summarized in Table 1. The performance of each model was measured using the mean average precision (mAP) metric at the intersection over union (IOU) (Rahman and Wang 2016) threshold of 0.5 and 0.75. To maintain consistency and ensure efficient processing, all images were resized to a standardized 512×512 pixel format. No additional image augmentation techniques were employed during the training process.

Name	Batch Size	Epochs	mAP^{50}	mAP^{75}	Time (hours)
d2-10k	8	$10,000$	43.56%	5.30%	9.52
d2-50k	8	$50,000$	80.60%	51.16%	12.22
d2-90k	8	$90,000$	82.65%	46.96%	22.50

Table 1: Detectron 2 training configurations

From this experiment, we found that *d2-50k* offered the optimal result offering the best performance while minimizing the training time.

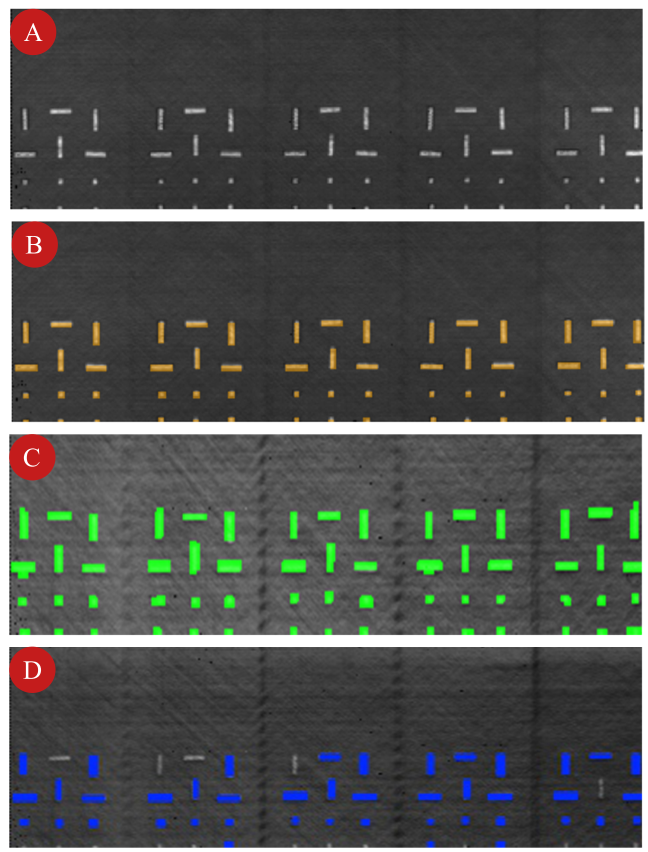

Figure 2: Exported images of the ultrasonic scans. A) The exported image used for the training. B) Orange labels showing the defect location on the image. C) Defects identified by the Detectron 2 model displayed with green labels. D) Defects identified by the YOLO 11 model displayed with blue labels.

Results. In our experiment, we achieved an average precision exceeding 80% at the 50% IOU threshold. However, given the critical safety implications of NDI applications, prioritizing sensitivity over precision is crucial. In this context, false positives are preferable to false negatives, as identifying every potential defect is essential. Since the primary goal of the model is to assist inspectors in locating regions of interest, detecting a fragment of a defect is sufficient to warrant a manual inspection of the affected area.

Figure 2 (C) visually represents the ability of the model to accurately identify all potential defect areas within a scan, highlighted by green overlays.

YOLO v11

We used the object detection models of YOLO 11, which integrates a backbone network for extracting features, followed by a segmentation head that generates both bounding boxes and detailed pixel-level masks for individual objects.

Training Configuration. To train the YOLO 11 model, we converted the COCO formatted data into appropriate labels and image files accepted by the YOLO (Ultralytics 2023) framework. We fine-tuned different hyperparameters of the pre-trained model to find the optimal configuration. We performed the training on the same hardware configuration as the Detectron 2 experiment. The performance of the training configurations is summarized in Table 2. To ensure efficient processing, all images were resized to standardized 640×640 pixels during training. We experimented with two pre-trained YOLO 11 models – *yolo11n*, a smallest model with $2.6 million$ parameters, and *yolo11x*, a larger model with $56.9 million$ parameters (Ultralytics 2024). The training processes were terminated when the model performance converged. The first experiment with $1,000$ epochs converged at 837 epochs whereas the other experiments converged at 1296 epochs.

Name	Batch Size	Max Epochs	mAP^{50}	mAP^{75}	Time (min)
yolo11n-1k	8	$1,000$	77.04%	48.41%	16.75
yolo11n-10k	8	$10,000$	75.22%	48.66%	29.55
yolo11x-10k	8	$10,000$	75.22%	48.66%	28.68

Table 2: YOLO 11 training configurations

Results. From the YOLO 11 training, we observed that the smaller pre-trained model, *yolo11n*, was sufficient to achieve optimal performance with minimal training time. Similar to the Detectron 2 experiment, the performance metric, while not perfect, was sufficient to identify almost all the defect areas, as illustrated by the blue overlays in Figure 2 (D).

Discussion

Our experiments demonstrated that the proposed method significantly reduces data pre-processing requirements. By converting scan data to PNG format, we successfully trained image segmentation models without requiring any complex pre-processing and data normalization. This streamlined pipeline enhanced the practicality of the models for industrial deployment by reducing computational overhead and increasing the model's adaptability across different scanning configurations.

The minimized pre-processing time enables real-time defect detection, allowing inspectors to identify defects as soon as the scanned data is available. This boosts efficiency and strengthens the NDI process. This minimized pre-processing step increased the model's viability across varied industrial environments.

Our experimental results indicate that the YOLO 11 model exhibited significantly faster training times while achieving performance comparable to the Detectron 2 model. This accelerated training process is particularly advantageous for rapid prototyping and iterative model development.

Compared to traditional ML models, which often require extensive signal preprocessing and cleanup, image segmentation models demonstrated remarkable performance directly from the raw scan data. This reduced reliance on complex pre-processing techniques simplifies the overall training and deployment pipeline for different manufactured materials.

Additionally, the high detection accuracy and efficiency of these models suggest their potential for practical application in real-world scenarios. The ability to accurately segment objects within images can be leveraged in a wide range of NDI processes.

Since the models rely on PNG images rather than underlying signals, they can easily adapt to defect detection on any new scan from a similarly shaped object. This flexibility minimizes the need for extensive retraining when introducing new manufactured parts. The retraining for new parts can also be simplified by creating images of the new components and incorporating them into the existing training data. This incremental approach allows for efficient adaptation to evolving product designs without requiring a complete retraining process.

Conclusion

Our experimental results present a robust approach to defect identification within the NDI process, leveraging state-of-the-art image segmentation techniques. The proposed methods demonstrate a high defect detection rate while minimizing the need for extensive data pre-processing and computational resources. The rapid detection capabilities inherent to these methods empower inspectors in the NDI process and significantly enhance productivity.

By integrating this approach into the defect detection pipeline, we can effectively reduce inspector fatigue, optimize productivity, and elevate the overall safety standards for manufactured parts. This integration can lead to a more efficient and reliable NDI process, reducing the likelihood of human error and improving the quality of the final product.

Acknowledgement

This work was supported in part by the U.S. Department of Energy, Office of Science, under contract **DE-AC02-06CH11357**.

References

Baumgartl, H.; Tomas, J.; Buettner, R.; and Merkel, M. 2020. A deep learning-based model for defect detection in laser-powder bed fusion using in-situ thermographic monitoring. *Progress in Additive Manufacturing*, 5(3): 277–285.

Brosnan, T.; and Sun, D.-W. 2004. Improving quality inspection of food products by computer vision—-a review. *J. Food Eng.*, 61(1): 3–16.

Cheng, B.; Collins, M. D.; Zhu, Y.; Liu, T.; Huang, T. S.; Adam, H.; and Chen, L.-C. 2020. Panoptic-DeepLab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

Diamanti, K.; Soutis, C.; and Hodgkinson, J. M. 2005. Non-destructive inspection of sandwich and repaired composite laminated structures. *Compos. Sci. Technol.*, 65(13): 2059–2067.

Dwivedi, S. K.; Vishwakarma, M.; and Soni, P. A. 2018. Advances and Researches on Non Destructive Testing: A Review. *Materials Today: Proceedings*, 5(2, Part 1): 3690–3698.

Elharrouss, O.; Al-Maadeed, S.; Subramanian, N.; Ottakath, N.; Almaadeed, N.; and Himeur, Y. 2021. Panoptic Segmentation: A Review.

Gardner, P.; Fuentes, R.; Dervilis, N.; Mineo, C.; Pierce, S. G.; Cross, E. J.; and Worden, K. 2020. Machine learning at the interface of structural health monitoring and non-destructive evaluation. *Philos. Trans. A Math. Phys. Eng. Sci.*, 378(2182): 20190581.

Gholizadeh, S. 2016. A review of non-destructive testing methods of composite materials. *Procedia Structural Integrity*, 1: 50–57.

Gupta, M.; Khan, M. A.; Butola, R.; and Singari, R. M. 2022. Advances in applications of Non-Destructive Testing (NDT): A review. *Advances in Materials and Processing Technologies*, 8(2): 2286–2307.

Hafiz, A. M.; and Bhat, G. M. 2020. A survey on instance segmentation: state of the art. *Int. J. Multimed. Inf. Retr.*, 9(3): 171–189.

Hassani, S.; and Dackermann, U. 2023. A Systematic Review of Advanced Sensor Technologies for Non-Destructive Testing and Structural Health Monitoring. *Sensors*, 23(4).

Honarvar, F.; and Varvani-Farahani, A. 2020. A review of ultrasonic testing applications in additive manufacturing: Defect evaluation, material characterization, and process control. *Ultrasonics*, 108: 106227.

Irfan, R.; Almazroi, A. A.; Rauf, H. T.; Damaševičius, R.; Nasr, E. A.; and Abdelgawad, A. E. 2021. Dilated Semantic Segmentation for Breast Ultrasonic Lesion Detection Using Parallel Feature Fusion. *Diagnostics (Basel)*, 11(7).

Jocher, G.; Jing, Q.; and Chaurasia, A. 2023. Ultralytics YOLO.

Jodhani, J.; Handa, A.; Gautam, A.; Ashwni; and Rana, R. 2023. Ultrasonic non-destructive evaluation of composites: A review. *Materials Today: Proceedings*, 78: 627–632.

Khanam, R.; and Hussain, M. 2024. YOLOv11: An overview of the key architectural enhancements. *arXiv [cs.CV]*.

Khedmatgozar Dolati, S. S.; Caluk, N.; Mehrabi, A.; and Khedmatgozar Dolati, S. S. 2021. Non-Destructive Testing Applications for Steel Bridges. *NATO Adv. Sci. Inst. Ser. E Appl. Sci.*, 11(20): 9757.

Kirillov, A.; He, K.; Girshick, R.; Rother, C.; and Dollar, P. 2019. Panoptic Segmentation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

Lange, Y. V. 1994. THE MECHANICAL IMPEDANCE ANALYSIS METHOD OF NONDESTRUCTIVE TESTING (A REVIEW). *NDT E Int.*, 11(4): 177–193.

Li, H.; Wang, C.; and Liu, Y. 2024. YOLO-FDD: efficient defect detection network of aircraft skin fastener. *Signal, Image and Video Processing*, 18(4): 3197–3211.

Lin, T.-Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Lawrence Zitnick, C.; and Dollár, P. 2014. Microsoft COCO: Common Objects in Context.

Lin, Y.; Ma, J.; Wang, Q.; and Sun, D.-W. 2023. Applications of machine learning techniques for enhancing nondestructive food quality and safety detection. *Crit. Rev. Food Sci. Nutr.*, 63(12): 1649–1669.

Ma, T.; Tsuchikawa, S.; and Inagaki, T. 2020. Rapid and non-destructive seed viability prediction using near-infrared hyperspectral imaging coupled with a deep learning approach. *Comput. Electron. Agric.*, 177: 105683.

Memmolo, P.; Arena, G.; Fatigati, G.; Grilli, M.; Paturzo, M.; Pezzati, L.; and Ferraro, P. 2015. Automatic Frames Extraction and Visualization From Noisy Fringe Sequences for Data Recovering in a Portable Digital Speckle Pattern Interferometer for NDI. *J. Display Technol.*, 11(5): 417–422.

Mishra, M.; Bhatia, A. S.; and Maity, D. 2020. Predicting the compressive strength of unreinforced brick masonry using machine learning techniques validated on a case study of a museum through nondestructive testing. *Journal of Civil Structural Health Monitoring*, 10(3): 389–403.

Ould Naffa, S.; Goueygou, M.; Piwakowski, B.; and Buyle-Bodin, F. 2002. Detection of chemical damage in concrete using ultrasound. *Ultrasonics*, 40(1-8): 247–251.

Prakash, N.; Nieberl, D.; Mayer, M.; and Schuster, A. 2023. Learning defects from aircraft NDT data. *NDT & E International*, 138: 102885.

Rahman, M. A.; and Wang, Y. 2016. Optimizing intersection-over-union in deep neural networks for image segmentation. In *Advances in Visual Computing*, Lecture notes in computer science, 234–244. Cham: Springer International Publishing. ISBN 9783319508344,9783319508351.

Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 779–788. IEEE. ISBN 9781467388511,9781467388528.

Su, Y.; Wang, Y.; Jiao, J.; and Guo, Y. 2011. Automatic detection and classification of breast tumors in ultrasonic images using texture and morphological features. *Open Med. Inform. J.*, 5(Suppl 1): 26–37.

Sun, H.; Ramuhalli, P.; and Jacob, R. E. 2023. Machine learning for ultrasonic nondestructive examination of welding defects: A systematic review. *Ultrasonics*, 127: 106854.

Taheri, H.; Gonzalez Bocanegra, M.; and Taheri, M. 2022. Artificial Intelligence, Machine Learning and Smart Technologies for Nondestructive Evaluation. *Sensors*, 22(11).

Thomas, J.; Kuang, B.; Wang, Y.; Barnes, S.; and Jenkins, K. 2024. Advanced semantic segmentation of aircraft main components based on transfer learning and data-driven approach. *The Visual Computer*, 1–20.

Ultralytics. 2023. Object Detection Datasets Overview. https://docs.ultralytics.com/datasets/detect. Accessed: 2024-11-21.

Ultralytics. 2024. YOLO11 NEW. https://docs.ultralytics.com/models/yolo11. Accessed: 2024-11-21.

Wang, B.; Zhong, S.; Lee, T.-L.; Fancey, K. S.; and Mi, J. 2020. Non-destructive testing and evaluation of composite materials/structures: A state-of-the-art review. *Advances in Mechanical Engineering*, 12(4): 1687814020913761.

Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019a. Detectron2. https://github.com/facebookresearch/detectron2.

Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019b. Detectron2.

Xiang, B.; Chang, H.; and Ye, Y. 2024. Aircraft Part Recognition from Drones' Perspective Based on Semantic Segmentation. In *2024 IEEE 7th International Conference on Electronic Information and Communication Technology (ICEICT)*, 1288–1292. IEEE.

Xiong, Y.; Liao, R.; Zhao, H.; Hu, R.; Bai, M.; Yumer, E.; and Urtasun, R. 2019. UPSNet: A Unified Panoptic Segmentation Network. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

Yang, Z.; Gao, W.; Chen, L.; Yuan, C.; Chen, Q.; and Kong, Q. 2022. A novel electromechanical impedance-based method for non-destructive evaluation of concrete fiber content. *Construction and Building Materials*, 351: 128972.